



Leveraging healthy population variability in deep learning unsupervised anomaly detection in brain FDG PET

Maëlys Solal, Ravi Hassanaly, Ninon Burgos

► To cite this version:

Maëlys Solal, Ravi Hassanaly, Ninon Burgos. Leveraging healthy population variability in deep learning unsupervised anomaly detection in brain FDG PET. SPIE Medical Imaging, Feb 2024, San Diego (California), United States. hal-04291561v2

HAL Id: hal-04291561

<https://hal.science/hal-04291561v2>

Submitted on 1 Feb 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Leveraging healthy population variability in deep learning unsupervised anomaly detection in brain FDG PET

Maëlys Solal^a, Ravi Hassanaly^a, and Ninon Burgos^a

^aSorbonne Université, Institut du Cerveau - Paris Brain Institute - ICM, CNRS, Inria, Inserm, AP-HP, Hôpital de la Pitié Salpêtrière, F-75013, Paris, France

ABSTRACT

Unsupervised anomaly detection is a popular approach for the analysis of neuroimaging data as it allows identifying a wide variety of anomalies from unlabelled data. It relies on reconstructing a subject-specific model of healthy appearance to which a subject's image can be compared to detect anomalies. In the literature, it is common for anomaly detection to rely on analysing the residual image between the subject's real image and its pseudo-healthy reconstruction. This approach however has limitations partly due to the pseudo-healthy reconstructions being imperfect and to the lack of natural thresholding mechanism. Our proposed method, inspired by Z-scores, leverages the healthy population variability to overcome these limitations. Our experiments conducted on 3D FDG PET scans from the ADNI database demonstrate the effectiveness of our approach in accurately identifying simulated Alzheimer's disease related anomalies.

Keywords: Unsupervised anomaly detection, Deep generative models, Alzheimer's disease, Abnormality maps

1. INTRODUCTION

Alzheimer's disease (AD) is a prevalent neurodegenerative disorder characterised by progressive memory loss and cognitive decline. Early detection of AD-related brain abnormalities is crucial, and neuroimaging such as ¹⁸F-fluorodeoxyglucose (FDG) positron emission tomography (PET) has the potential to play a vital role in this process, by allowing to visualise subtle changes in neuronal activity before the onset of clinical symptoms.¹

Machine learning (ML) methods have shown promise in the quantitative analysis of neuroimaging data, which is key in detecting subtle changes in brain patterns amidst natural variations. This work aligns with the unsupervised anomaly detection (UAD) framework, where ML algorithms are used to identify rare events from unlabelled data.² The use of UAD algorithms is motivated by the limited availability of labelled data, and the hope that these methods enable the detection of various types of anomalies, including those associated with AD.

This approach relies on generating a subject-specific model of healthy appearance, and comparing the subject's real image with the model, resulting in a subject-specific map of anomalies. In the training phase, the neural network is exclusively shown images of healthy subjects, enabling it to learn the distribution of healthy images. When applied to a patient's image (with an unknown diagnosis) during the testing phase, the network reconstructs the image based on its knowledge of healthy images.² Consequently, we hypothesise that the reconstructed image is a subject-specific pseudo-healthy representation of the input image. By comparing the input image with its reconstruction, areas of the brain exhibiting anomalies can be effectively highlighted.

One popular method in anomaly detection involves analysing the residual error between a patient's image and its pseudo-healthy reconstruction.^{2,3} This approach however has several shortcomings.⁴ Firstly, there is no clear thresholding mechanism to selectively focus on significant anomalies, which leads to a higher number of false positives. Secondly, the method does not consider the normal variability within the dataset, making it harder to capture subtle anomalies. Additionally, it fails to account for the reconstruction error of the model itself, which may be evident in images from healthy individuals.

Further author information: (Send correspondence to Maëlys Solal)

Maëlys Solal: E-mail: maelys.solal@icm-institute.org

Ninon Burgos: E-mail: ninon.burgos@cnrs.fr

Our proposed approach aims to address the aforementioned limitations by introducing abnormality maps constructed by relying loosely on Z-scores. Our method leverages variability within the training dataset to enhance anomaly detection. It also provides improved interpretability, as it allows for a direct association between the abnormality map and the standard deviation within a control population, thereby facilitating the identification of significant anomalies by establishing a clear threshold. Consequently, our approach offers a more comprehensive solution to design abnormality maps that better account for model errors, and are more interpretable, by leveraging the intuitive nature of Z-scores.

2. MATERIALS AND METHODS

2.1 Dataset

In this study, we work with FDG PET scans from the ADNI database.⁵⁻⁷ We select images that are co-registered, averaged and uniformised to the same resolution. These images are then further pre-processed using the `pet-linear` pipeline from Clinica⁸ which registers to a standard space, normalises in intensity and crops each image. For our purposes, we work with images from healthy subjects, labelled as cognitively normal (CN), and some images of diseased subjects, labelled as Alzheimer’s disease (AD). In total, we include 318 CN subjects (which amounts to 590 images) and 273 AD subjects, with only one image each.

2.2 Generative model - 3D VAE

To generate subject-specific models of healthy appearance, we rely on variational autoencoders (VAEs),⁹ a type of deep probabilistic generative model. During the training phase, the model is shown only images labelled as CN to learn the distribution of healthy images. At inference time, the image of a subject is given as input to the previously trained VAE, which generates an image that looks healthy and corresponds to that subject specifically. By comparing the input image and its pseudo-healthy reconstruction outputted by the VAE, we can detect and localise the areas of the brain which present anomalies.

A VAE aims to approximate the true distribution of the training data with a simple parameterised distribution. In practice, the encoder projects an input sample \mathbf{x} to a latent space of smaller dimension which is constrained to parameterise a standard normal distribution. Each input is thus mapped to a normal distribution $\mathcal{N}(\boldsymbol{\mu}(\mathbf{x}), \boldsymbol{\sigma}(\mathbf{x})^2 \mathbf{I})$. The latent vector \mathbf{z} is then sampled from this distribution, and the decoder learns to reconstruct $\hat{\mathbf{x}}$ from \mathbf{z} .

2.3 Abnormality maps

Once the subject-specific pseudo-healthy image $\hat{\mathbf{x}}$ is generated, it needs to be compared to the input image \mathbf{x} to generate an abnormality map from which anomalies can be detected. In the literature, it is common to define the abnormality map as the voxel-wise difference between these two images $\mathbf{r} = \mathbf{x} - \hat{\mathbf{x}}$, also known as the residual image. Unfortunately, this method does not account for the normal variability in the population and the potential reconstruction errors of our generative model. In other words, the residual errors are considered at the same level of importance for every voxel, even in regions which are highly variable between subjects.

In a typical voxel-wise analysis,^{10,11} a subject’s PET image is registered to a standard space and compared voxel-by-voxel to a distribution obtained from normal control scans. This can be done through a Z-score, or standard score, given by the ratio of the difference between an individual raw value and the control population mean, and the control population standard deviation. The score then corresponds to the number of standard deviations by which the observed value deviates from the mean value in a control population. Classically, this method also allows setting a threshold, corresponding to a level of confidence.

In our work, we generate abnormality maps which rely loosely on the idea of Z-scores since this statistic is widely used among clinicians. We set: $\tilde{\mathbf{z}} = \frac{\mathbf{x} - \hat{\mathbf{x}}}{\sigma}$, where \mathbf{x} is the input image, $\hat{\mathbf{x}}$ is its pseudo-healthy reconstruction, and σ is the standard deviation of the set of healthy images used to train our model (in some way, our healthy control population).

2.4 AD hypometabolism simulation framework

Evaluating our ability to detect anomalies is challenging due to the lack of ground truth. Indeed, we only have access to unlabelled images, with a diagnosis label but no anomaly annotations, that is, neither contour nor mask. We can however rely on the hypometabolism simulation framework proposed by Ref. 12,13. It consists of a new test set, where hypometabolism characteristic of Alzheimer’s disease (or other dementia subtypes) can be simulated on healthy images to have pairs of healthy and diseased images with known anomalies. For a given healthy image \mathbf{x} , we obtain a corresponding image \mathbf{x}' with a certain degree of hypometabolism (in our case, 30%). To validate our ability to detect anomalies, we can compare our abnormality map (residual-based or Z-score-based) with the mask that was used to simulate the hypometabolism between \mathbf{x} and \mathbf{x}' in the first place.

2.5 Experimental setup

To avoid data leakage, we split our dataset into distinct splits. The images of the AD subjects are used solely during testing. We perform training, validation and testing data splits of the CN dataset at the subject’s level and stratified by sex and age to reduce biases. There are three different test tests, each with a single image from each subject: (i) 60 images from CN subjects, (ii) 60 images with simulated hypometabolism, (iii) and 273 images from AD subjects. The validation set includes 53 images from 53 CN subjects, and the training set includes 537 images from 265 CN subjects.

To synthesise images of healthy appearance, we rely on the 3D VAE proposed in Ref. 13, which we represent in Fig. 1. Its encoder is composed of five convolutional layers (where the size of the image is divided by 2 along each dimension) and one dense layer, and its decoder is symmetrical, by adding an upsampling layer before each convolution. After each convolutional layer is a batch normalisation and a swish activation function, except for the last layer, where there is no batch normalisation, and the activation function is a sigmoid. The latent space was chosen to have a size of 256. The model was trained on 200 epochs, with a learning rate of 10^{-4} and a batch size of 8 using ClinicaDL.¹⁴




Figure 1: Variational autoencoder (VAE) architecture

3. RESULTS

To assess the inter- and intra-subject variability of our data, we performed both a voxel-wise and a regional statistical analysis, and report the results of the latter in Fig. 2. Both analyses confirm our intuition that some regions exhibit a larger variability than others, making the voxels in that region more difficult to reconstruct, and therefore more prone to errors when performing anomaly detection. It therefore seems coherent to attempt to account for this variability in our abnormality maps. For instance, the threshold to consider a voxel as abnormal in regions that exhibit high inter-subject variability could be set higher than in those with lower variability.




Figure 2: Preliminary analysis: comparison of the distribution of average regional FDG uptake in cognitively normal (orange) and Alzheimer’s disease (blue) subjects

We plot the results of residual-based and Z-score-based abnormality maps for a subject with simulated hypometabolism in Fig. 3. We find that Z-score-based abnormality maps indeed better highlight relevant anomalies simulated with the hypometabolism mask through the use of thresholds, thus reducing the impact of the model’s reconstruction error.

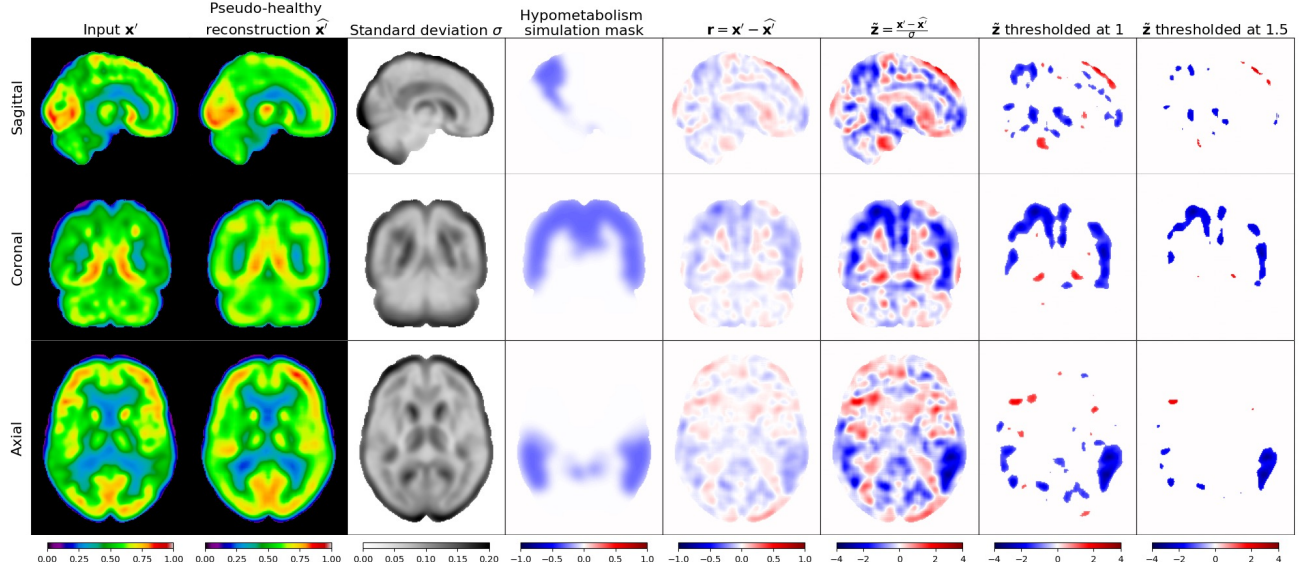
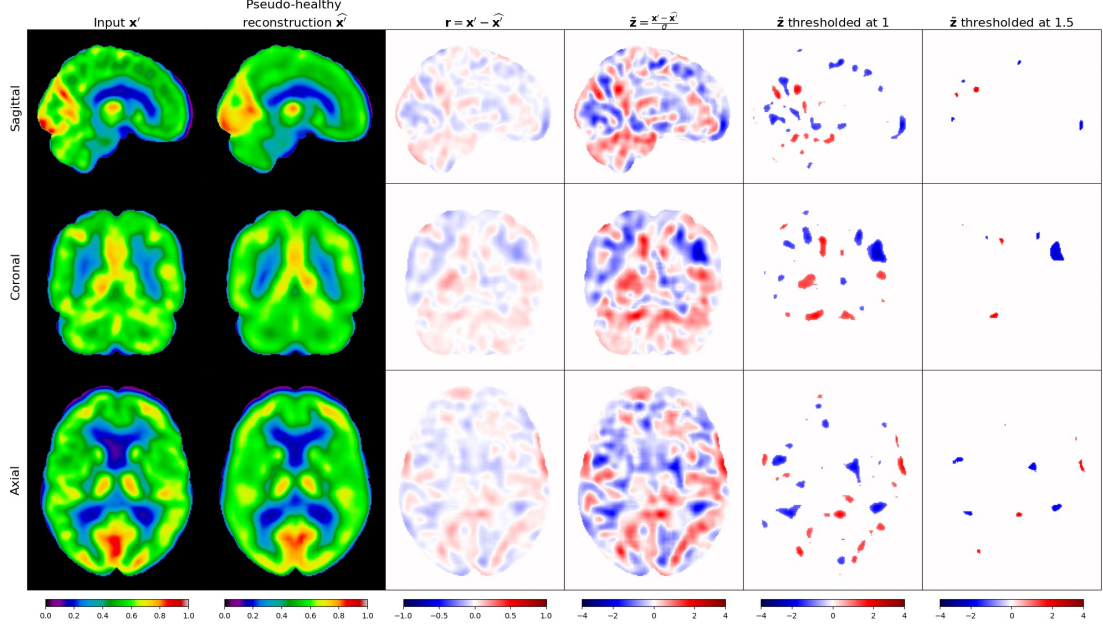


Figure 3: Example of results obtained for a healthy image with simulated hypometabolism. We display a central slice for each plane with the input \mathbf{x} , the pseudo-healthy reconstruction $\hat{\mathbf{x}}$, the standard deviation of the training set of CN images σ , the hypometabolism simulation mask with the anomalies we aim to retrieve, the residual $\mathbf{r} = \mathbf{x} - \hat{\mathbf{x}}$, the abnormality map $\hat{\mathbf{z}} = \frac{\mathbf{x} - \hat{\mathbf{x}}}{\sigma}$, and the abnormality map $\hat{\mathbf{z}}$ thresholded at 1 and 1.5.

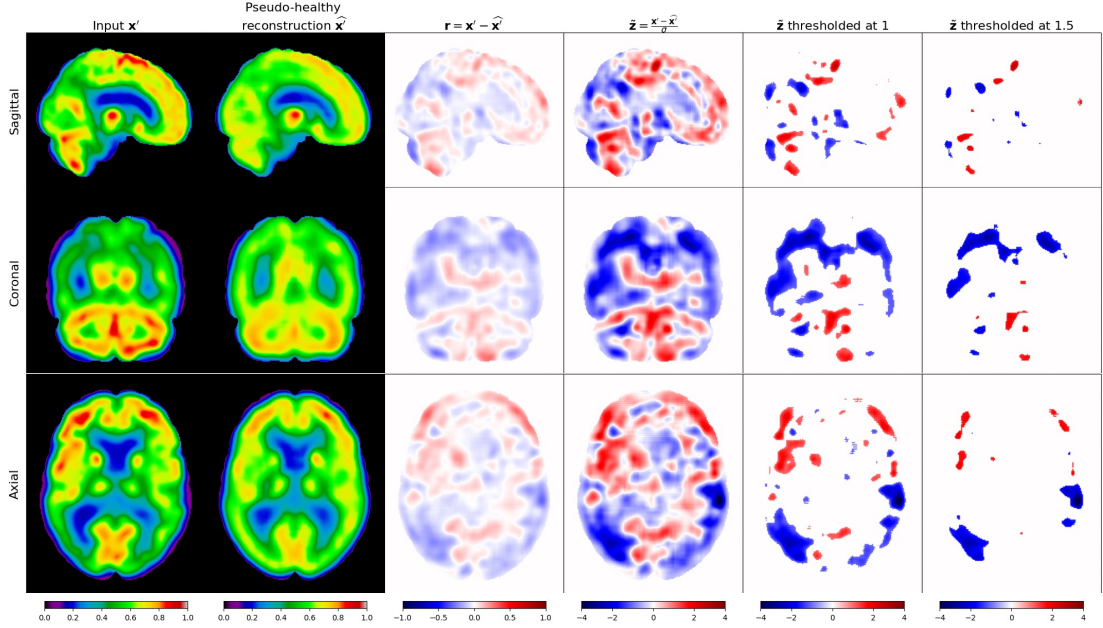
To complement our visual analysis of the abnormality maps, we compute the normalised cross-correlation coefficient (NCC) between the hypometabolism simulation mask and our abnormality maps. The average NCC across the testing set between the mask and our proposed abnormality map $\hat{\mathbf{z}}$ is $0.568 (\pm 0.059)$, which is significantly higher than the average NCC between the mask and the residual \mathbf{r} is $0.510 (\pm 0.060)$ ($p < 10^{-4}$, computed with a t-test). We find that the Z-score-based abnormality maps allow us to better retrieve the anomalies that were simulated, as shown by a larger correlation magnitude.

For the sake of comparison, we also include results from a cognitively normal subject and an Alzheimer’s subject (Fig. 4). Unfortunately, due to the lack of ground truth and appropriate metrics, it is difficult to display

and compare results for the entire testing sets. While there are remaining artefacts in the abnormality maps for the CN subject, we can see that this approach highlights fewer anomalies in comparison to the residual approach, making it more robust to noise. As for the AD subject, we can see that the anomalies are well highlighted in places where we expect to find them, such as the temporal lobe.



(a) Cognitively normal subject



(b) Alzheimer's disease patient

Figure 4: Example of results obtained for images from real CN and AD subjects. We display a central slice for each plane with the input \mathbf{x} , the pseudo-healthy reconstruction $\hat{\mathbf{x}}$, the residual $\mathbf{r} = \mathbf{x} - \hat{\mathbf{x}}$, the abnormality map $\hat{\mathbf{z}} = \frac{\mathbf{x} - \hat{\mathbf{x}}}{\sigma}$, and the abnormality map $\hat{\mathbf{z}}$ thresholded at 1 and 1.5.

4. DISCUSSION & CONCLUSION

Our work explores the detection of AD-related brain abnormalities using FDG PET neuroimaging. We adopt a UAD approach which relies on building subject-specific pseudo-healthy neuroimages to which a patient’s image can be compared, as it provides the ability to detect subtle changes without prior knowledge. Traditional methods rely on residual analysis which has limitations which include the lack of natural thresholding mechanism, and its failure to consider normal variability and reconstruction errors.²⁻⁴

To address these limitations, we propose a novel approach that leverages healthy population variability to construct abnormality maps. Our method is inspired from Z-scores and offers improved interpretability by establishing a clearer association between the anomalies and the standard deviation in a control population. The integration of Z-scores also facilitates the identification of significant anomalies by providing a more natural thresholding mechanism. Our experimental results on FDG PET scans from the ADNI database highlight the potential of our method in accurately identifying simulated brain abnormalities related to AD. Future works could involve validating this approach with other types of anomalies, such as those linked with other types of dementia, and further evaluating this method with images from real AD subjects.

Whilst our method allows building more interpretable abnormality maps that highlight significant anomalies, it remains challenging to distinguish reconstruction errors caused by the model being imperfect from those related to anomalies in the input. Future efforts could be directed towards improving the reconstruction capabilities of the generative model, especially since VAEs are known for producing blurry reconstructions,² with diffusion models¹⁵ or conditional GANs for instance.¹⁶ We could also try post-processing techniques to mitigate the effects of reconstruction errors, or generating abnormality maps using larger regions instead of voxel-wise analysis.

In summary, our proposed approach combines the advantages of UAD, such as the ability to detect subtle changes without prior knowledge, with the interpretability offered by Z-scores. Our experiments provide promising results, indicating that our method holds the potential for more reliable early detection of AD-related brain abnormalities.

ACKNOWLEDGMENTS

The research leading to these results has received funding from the French government under management of Agence Nationale de la Recherche as part of the "Investissements d’avenir" program, reference ANR-19-P3IA-0001 (PRAIRIE 3IA Institute) and reference ANR-10-IAIHU-06 (Agence Nationale de la Recherche-10-IA Institut Hospitalo-Universitaire-6).

The ADNI was launched in 2003 as a public-private partnership, led by Principal Investigator Michael W. Weiner, MD. The primary goal of ADNI has been to test whether serial MRI, PET, other biological markers, and clinical and neuropsychological assessment can be combined to measure the progression of mild cognitive impairment and early AD.

REFERENCES

- [1] Mosconi, L., “Brain glucose metabolism in the early and specific diagnosis of Alzheimer’s disease: FDG-PET studies in MCI and AD,” *European Journal of Nuclear Medicine and Molecular Imaging* **32**(4), 486–510 (2005).
- [2] Chen, X. and Konukoglu, E., “Unsupervised abnormality detection in medical images with deep generative methods,” in *[Biomedical Image Synthesis and Simulation]*, Burgos, N. and Svoboda, D., eds., 303–324, Elsevier (2022).
- [3] Baur, C., Denner, S., Wiestler, B., Navab, N., and Albarqouni, S., “Autoencoders for unsupervised anomaly segmentation in brain MR images: A comparative study,” *Medical Image Analysis* **69**, 101952 (2021).
- [4] Meissen, F., Wiestler, B., Kaissis, G., and Rueckert, D., “On the Pitfalls of Using the Residual as Anomaly Score,” in *[Medical Imaging with Deep Learning]*, (2021).
- [5] Mueller, S. G., Weiner, M. W., Thal, L. J., Petersen, R. C., Jack, C. R., Jagust, W., Trojanowski, J. Q., Toga, A. W., and Beckett, L., “Ways toward an early diagnosis in Alzheimer’s disease: The Alzheimer’s Disease Neuroimaging Initiative (ADNI),” *Alzheimer’s & Dementia* **1**(1), 55–66 (2005).

- [6] Jagust, W. J., Bandy, D., Chen, K., Foster, N. L., Landau, S. M., Mathis, C. A., Price, J. C., Reiman, E. M., Skovronsky, D., and Koeppe, R. A., “The Alzheimer’s Disease Neuroimaging Initiative Positron Emission Tomography Core,” *Alzheimer’s & Dementia* **6**(3), 221–229 (2010).
- [7] Jagust, W. J., Landau, S. M., Koeppe, R. A., Reiman, E. M., Chen, K., Mathis, C. A., Price, J. C., Foster, N. L., and Wang, A. Y., “The Alzheimer’s Disease Neuroimaging Initiative 2 PET Core: 2015,” *Alzheimer’s & Dementia* **11**(7), 757–771 (2015).
- [8] Routier, A., Burgos, N., Díaz, M., Bacci, M., Bottani, S., El-Rifai, O., Fontanella, S., Gori, P., Guillon, J., Guyot, A., Hassanaly, R., Jacquemont, T., Lu, P., Marcoux, A., Moreau, T., Samper-González, J., Teichmann, M., Thibeau-Sutre, E., Vaillant, G., Wen, J., Wild, A., Habert, M.-O., Durrleman, S., and Colliot, O., “Clinica: An Open-Source Software Platform for Reproducible Clinical Neuroscience Studies,” *Frontiers in Neuroinformatics* **15**, 39 (2021).
- [9] Kingma, D. P. and Welling, M., “Auto-Encoding Variational Bayes,” (2014).
- [10] Minoshima, S., Frey, K., Koeppe, R., Foster, N., and Kuhl, D., “A diagnostic approach in Alzheimer’s disease using three-dimensional stereotactic surface projections of Fluorine-18-FDG PET,” *Journal of Nuclear Medicine* **36**, 1238–48 (1995).
- [11] Drzezga, A., Grimmer, T., Riemenschneider, M., Lautenschlager, N., Siebner, H., Alexopoulos, P., Minoshima, S., Schwaiger, M., and Kurz, A., “Prediction of individual clinical outcome in MCI by means of genetic assessment and 18F-FDG PET,” *Journal of Nuclear Medicine* **46**(10), 1625–1632 (2005).
- [12] Hassanaly, R., Bottani, S., Sauty, B., Colliot, O., and Burgos, N., “Simulation based evaluation framework for deep learning unsupervised anomaly detection on brain FDG-PET,” in [*Medical Imaging 2023: Image Processing*], **12464**, 1246428, SPIE (2023).
- [13] Hassanaly, R., Brianceau, C., Solal, M., Colliot, O., and Burgos, N., “Evaluation of pseudo-healthy image reconstruction for anomaly detection with deep generative models: Application to brain FDG PET,” *Machine Learning for Biomedical Imaging* **2**(Special Issue for Generative Models), 611–656 (2024).
- [14] Thibeau-Sutre, E., Díaz, M., Hassanaly, R., Routier, A., Dormont, D., Colliot, O., and Burgos, N., “ClinicaDL: An open-source deep learning software for reproducible neuroimaging processing,” *Computer Methods and Programs in Biomedicine* **220** (2022).
- [15] Pandey, K., Mukherjee, A., Rai, P., and Kumar, A., “Diffusevae: Efficient, controllable and high-fidelity generation from low-dimensional latents,” (2022). arXiv:2201.00308.
- [16] Yaakub, S., McGinnity, C., Clough, J., Kerfoot, E., Girard, N., Guedj, E., and Hammers, A., “Pseudo-normal PET synthesis with generative adversarial networks for localising hypometabolism in epilepsies,” in [*Simulation and Synthesis in Medical Imaging*], *LNCS*, 42–51, Springer (2019).